# Semi-supervised Learning with a Teacher-student Network for Generalized Attribute Prediction

Minchul Shin<sup>1</sup>[0000-0002-0638-2017]</sup>

Search Solutions Inc., Republic of Korea min.stellastra@gmail.com

Abstract. This paper presents a study on semi-supervised learning to solve the visual attribute prediction problem. In many applications of vision algorithms, the precise recognition of visual attributes of objects is important but still challenging. This is because defining a class hierarchy of attributes is ambiguous, so training data inevitably suffer from class imbalance and label sparsity, leading to a lack of effective annotations. An intuitive solution is to find a method to effectively learn image representations by utilizing unlabeled images. With that in mind, we propose a multi-teacher-single-student (MTSS) approach inspired by the multi-task learning and the distillation of semi-supervised learning. Our MTSS learns task-specific domain experts called teacher networks using the label embedding technique and learns a unified model called a student network by forcing a model to mimic the distributions learned by domain experts. Our experiments demonstrate that our method not only achieves competitive performance on various benchmarks for fashion attribute prediction, but also improves robustness and cross-domain adaptability for unseen domains.

Keywords: Semi-supervised Learning, Unlabeled Data, Visual Attributes

# 1 Introduction

Visual attributes are qualities of objects that facilitate the cognitive processes of human beings. Therefore, predicting the attributes of an object accurately has many useful applications in the real world. For example, a search engine can use predictions to screen products with undesirable attributes instead of using noisy metadata provided by anonymous sellers [2]. Attribute prediction is essentially a multi-label classification problem that aims to determine if an image contains certain attributes (*e.g., colors and patterns*). However, attribute prediction is known as a very challenging task based on the expense of annotation, difficulty in defining a class hierarchy, and simultaneous appearance of multiple attributes in objects. Although recent works have shown competitive results on various benchmark datasets [7, 12], we identified several additional conditions that must be satisfied to solve the aforementioned issues for such methods to be useful in



(a) Multi-Instance, Single-Attribute (b) Single-Instance, Multi-Attribute

**Fig. 1.** Two different types of visual attribute prediction tasks. Image (a) was sampled from MSCOCO [21] and image (b) was sampled from DeepFashion [23].

real-world applications: domain-agnostic training, the use of unlabeled data, and robustness/generalization.

Use of unlabeled data. The human perception of an attribute is intrinsically subjective, leading to sparse and ambiguous annotations in datasets. Additionally, attributes have very-long-tailed distributions, leading to severe class imbalance (*e.g., pattern: solid versus zebra*). Therefore, even if the total number of annotated images is large, models can suffer from extreme shortages of usable training data following class balancing in many cases [9]. An intuitive solution is to use unlabeled images for training in a semi-supervised manner to enhance the ability of a model to represent the visual information included in images. The teacher-student paradigm proposed in this paper facilitates this process with no preprocessing tricks, such as obtaining pseudo-labels based on predictions [5, 41, 40].

**Robustness/generalization.** Robustness measures how stable and reliable a network is when making decisions in the presence of unexpected perturbations in inputs [43]. Generalization measures how well a trained model performs with a domain shift or if a model can be generalized sufficiently with insufficient training data [17, 26]. In addition to high accuracy on benchmarks, these aspects must also be considered as crucial factors for measuring model quality.

**Domain-agnostic training.** To achieve state-of-the-art performance on benchmarks, recent works [42, 23, 22, 4] have relied on domain-specific auxiliary information during training. However, such methods have obvious limitations in terms of expansion to additional domains, which is an important feature for real applications. Therefore, it is preferable to avoid using any domain-specific information during training. We refer to this principle as domain-agnostic training.

Main contributions. In this paper, we propose a multi-teacher-single-student (MTSS) method for generalized visual attribute prediction that aims to train a single unified model to predict all different attributes in a single forward operation. This paper makes the following main contributions. First, we introduce an MTSS method that learns from multiple domain experts in a semi-supervised



Fig. 2. Conceptual summary of previous methods and the proposed method denoted, as common [42, 23, 22, 4], progressive [2], and MTSS.

manner utilizing unlabeled images. We also show the advantages of label embedding for training teacher networks, which are used as domain experts. Second, we demonstrate that the MTSS method can enhance model quality in terms of robustness and generalization without sacrificing benchmark performance by learning from images in distinct domains. Finally, our approach is fully domain agnostic, meaning it requires no domain-specific auxiliary supervision, such as landmarks, pose detection, or text description, for training. By learning using only attribute labels, our model not only outperforms previous methods under the same conditions [8, 6, 16], but also achieves competitive results relative to existing domain-specific methods [23, 38, 32, 22] that use additional auxiliary labels for supervision.

## 2 Related Work

**Visual Attributes.** Visual attribute prediction (VAP) is a multi-label classification task that has been widely studied [42, 4, 2, 22, 38, 23, 32, 30]. There are two main types of VAP tasks, as illustrated in Figure 1. We are particularly interested in type (b), which is frequently observed in the fashion domain. This is because our interests lie in the ability of a model to predict multiple simultaneously appearing attributes, rather than localizing interest regions, which is a matter of detection. The attributes appearing in type (b) are not necessarily related to visual similarity. For example, they may only be related to low-level characteristics, such as color, which makes the task more interesting. In the fashion domain, some works [42, 23, 22, 4] have used landmarks for clothing items, pose detection, or textual item descriptions to improve overall accuracy. However, such strong requirements regarding auxiliary information limit such methods to domain-specific solutions. In contrast, some works have focused on attention mechanisms [38, 32, 42]. Zhang et al. [42] proposed a task-aware attention mechanism that considers the locality of clothing attributes according to different tasks. Wang et al. [38] proposed landmark-aware attention and category-driven attention. The former focuses on the functional parts of clothing items and the latter enhances task-related features that are learned in an implicit manner. VSAM [32] uses the inference results of pose detection for supervision to train a pose-guided attention model. Although attention mechanisms are well-known methods for boosting performance, such mechanisms mainly discuss where to focus given spatial regions. Because we consider attention mechanisms as an extension that can be applied to any existing method with slight modifications, such mechanisms are given little attention in this work. One of main challenges in VAP is solving multi-task learning problems in the presence of label sparsity and class imbalances in training data [1]. A single unified model is preferred for saving on the cost of inference and is expected to provide robust predictions for all targets (e.g., style, texture, and patterns in fashion). The most commonly used method is to train multiple binary classifiers [23,8]. However, because the numbers of effective annotations for each task differ significantly based on label sparsity, models can easily encounter overfitting on tasks with many annotations and underfitting on tasks with few annotations over the same number of iterations. Balancing such bias during training is important issue that must be solved for VAP. Adhikari et al. [2] proposed a progressive learning approach that add branches for individual models for attributes progressively as training proceeds. Lu et al. [24] proposed a tree-like deep architecture that automatically widens the network for multi-task learning in a greedy manner. However, the implementation of such methods is tricky and requires considerable engineering work.

Label Embedding. Label embedding (LE) refers to an important family of multi-label classification algorithms that have been introduced in various studies [3, 36, 28]. Such approaches jointly learn mapping functions  $\varphi$  and  $\lambda$  that project embeddings of an image x and label y in a common intermediate space Z. Compared to direct attribute prediction, which requires training a single binary classifier for each attribute, LE has the following advantages. First, attributes are not required to be independent because one can simply move the embedding of x closer to the correct label y than to any incorrect labels y' based on ranking loss [15, 37]. Second, prediction classes can be readily expanded because classes can be predicted by measuring the shortest distance to the center point of a feature cluster assigned to an unseen class [19, 35, 18, 33]. We found a number of studies examining LE in terms of zero-shot learning. In this paper, we highlight the usage of LE to train domain experts for an MTSS model for semi-supervised learning.

Semi-supervised Learning and Distillation. Recent works [5, 41, 40] have investigated semi-supervised learning (SSL) methods [10] that use large-scale unlabeled images to improve supervision. Yalniz *et. al.* [41] proposed a self-training method representing a special form of SSL and achieved state-of-the-art accuracy by leveraging a large amount of unannotated data. Distillation can also be considered as a form of SSL in that a teacher model makes predictions for unlabeled data and the results are used for supervision to train a student model [29, 11, 31]. Such strategies have yielded impressive results on many vision tasks. Inspired by these methods, we propose a training method that takes advantage of SSL for VAP. We found that a teacher-student paradigm of distillation is very effective for performing VAP because relevant visual attributes can be grouped into particular attribute types (*e.g.* {*stripe*, *dot*}  $\in$  *pattern*), meaning a student model can effectively learn from multiple teachers that are specialists for each attribute type.

# 3 Methodology

## 3.1 Overview

The main concept of the MTSS approach is to integrate multiple teachers (MT) that are domain experts for each attribute type into a single student (SS) to construct a unified model that can predict multiple attributes in a single forward operation. Given an image x and attribute type  $\alpha_k \in \{\alpha_1, \alpha_2, ..., \alpha_\kappa\}$ , our goal is to predict the attribute classes  $c_p \in \{c_1^{\alpha_k}, c_2^{\alpha_k}, ..., c_P^{\alpha_k}\}$ , where  $\kappa$  and P are the number of target attribute types and final predictions to be outputted, respectively. Suppose  $\alpha_k$  is a *pattern* and *color*. Then, the corresponding  $c_p$  could be dot, stripe or red. P may differ depending on which confidence score is used for the predicted result to reach a final decision regarding each attribute type  $\alpha_k$ . Our assumption is that attributes existing in the real world have a conceptual hierarchy, implying that relevant visual attributes are grouped into a particular attribute type  $\alpha_k$ . For example, > 1000 attributes in DeepFashion are grouped into 6 attribute types which means only 6 teachers are required for training student. The training procedure can be divided into two stages for the teacher and student. The design details and advantages of our two-stage training method are discussed in the following subsections.

#### 3.2 Teacher Models for Individual Attributes

Given pairs of an image x and ground-truth label y in a training set  $\zeta_{\alpha_k} = \{(x_n, y_n), n = 1, ..., N\}$  for an attribute type  $\alpha_k$  with  $x_n \in X$  and  $y_n \in Y$ , our goal is to train a teacher model  $\varphi_{\alpha_k}^T : X \to Y$ . Our teacher model uses the pairwise ranking loss of metric learning to learn image representations. Specifically, a group of randomly initialized label embeddings obtained from an attribute dictionary  $d_{\alpha_k}^n = \lambda_{\alpha_k}(y_n)$ , where  $d_{\alpha_k}^n \in R^D$  and D denotes the dimension of the image embedding, is defined and  $d_{\alpha_k}^n$  is learned to represent the center of the feature cluster of the *n*-th class for attribute type  $\alpha_k$ . Given a label embedding  $d_{\alpha_k}^n$  that corresponds to the most representative feature point of the *n*-th class of  $\alpha_k$ , our goal is to locate an embedding of the positive image  $e^+ = \varphi_{\alpha_k}^T(I^+)$  to be close to  $d_{\alpha_k}^n$  and an embedding of the negative image  $e^- = \varphi_{\alpha_k}^T(I^-)$  far away. The positive image is sampled from the same attribute class as  $d_{\alpha_k}^n$  and the negative image is sampled from a randomly selected class that is different from the class of the positive image. Although various sampling strategies can be adopted to enhance overall performance [39], such optimization was omitted in this work because it lies outside of our focus. The general form of an objective can be written as (1).

$$\ell'(d^n_{\alpha_k}, e^+, e^-) = max\{0, (1 - \|d^n_{\alpha_k}\|_2 \cdot \|e^+\|_2 + \|d^n_{\alpha_k}\|_2 \cdot \|e^-\|_2)\}, \qquad (1)$$



**Fig. 3.** Training of a teacher model. The feature extractor  $\varphi^T$  and mapping  $\lambda$  are learned to place the image embedding (left) and label embedding (right) close to each other in latent space.

where  $\|\|\|_2$  represents L2 normalization. Lin *et al.* [20] introduced the concept of focal loss, which can help a model focus on difficult and misclassified examples by multiplying a factor  $(1 - p_t)^{\gamma}$  by the standard cross-entropy loss, where  $p_t$  is the estimated probability for the target class. As a modification for metric learning, we applied focal loss to our method by defining the probability  $p_t$  as the cosine similarity between d and e, which is bounded in the range of [0, 1]. The final objective is defined as (2).

$$p_t = 0.5 \times max\{0, (1 + \|d^n_{\alpha_k}\|_2 \cdot \|e^+\|_2 - \|d^n_{\alpha_k}\|_2 \cdot \|e^-\|_2)\}, \\ \ell^{Teacher}(p_t) = -(1 - p_t)^{\gamma} log(p_t),$$

$$(2)$$

where  $\gamma$  is the hyperparameter to be found. Note that  $\gamma$  was set to 1.0 in our experiments unless stated otherwise. We found that if  $\gamma$  is adjusted carefully, adopting the modified focal loss term significantly boosts overall scores, which will be discussed in the experimental section. During training, an image can be duplicated in more than one attribute class if the multiple labels exist in the ground truth. We expect the number of experts (i.e., trained teacher models  $\varphi^T \in \{\varphi^T_{\alpha_1}, \varphi^T_{\alpha_2}, ..., \varphi^T_{\alpha_k}\}$ ) to be the same as the number of attribute types  $\kappa$ .

#### 3.3 Unified Student Model for All Attributes

The goal of the student stage is to integrate the trained teacher models  $\varphi^T \in \{\varphi_{\alpha_1}^T, \varphi_{\alpha_2}^T, ..., \varphi_{\alpha_k}^T\}$  into a single unified network  $\varphi^S$  and boost overall performance by utilizing unlabeled images  $\mathcal{U}$  in a semi-supervised manner. A distillation method that aims to transfer knowledge acquired in one model, namely a teacher, to another model, namely a student, was adopted for this purpose. The core idea of distillation is very simple. A student model is trained to mimic the feature distribution of  $\varphi^T$ . Learning distributions from trained models can be



Fig. 4. Training of a student model. The teachers, which represent individual experts for each attribute type, are independent models, unlike the unified student model that shares weights before its fully connected layers.

achieved simply by matching the image embedding of a teacher  $e_{\alpha_k}^T = \varphi_{\alpha_k}^T(x)$  to the embedding of a student  $e_{\alpha_k}^S = \varphi^S(x; \alpha_k)$  according to a target attribute type  $\alpha_k$ . It is assumed that if  $\varphi^S$  is able to reproduce the same feature distribution as  $\varphi^T$ , then prediction can be performed by simply measuring the distance between  $e_{\alpha_k}^S$  and a learned dictionary  $d_{\alpha_k} \in \{d_{\alpha_k}^1, d_{\alpha_k}^2, ..., d_{\alpha_k}^N\}$  that was already obtained in the teacher stage. A student  $\varphi^S$  consists of S branches of fully-connected layers following the last pooling layer of the backbone, where each branch is in charge of projecting pooled descriptors into an attribute-specific embedding space  $Z_{\alpha_k} \in \mathbb{R}^D$ . The weights before the fully connected branches are shared. Given an attribute type  $\alpha_k$ , we maximize the cosine-similarity between  $e_{\alpha_k}^S$  and  $e_{\alpha_k}^T$ . Therefore objective is formulated as

$$\ell'' = \Sigma_{k=1}^{\kappa} \{ 1 - \| e_{\alpha_k}^S \|_2 \cdot \| e_{\alpha_k}^T \|_2 \}, \tag{3}$$

where  $\|\|_2$  is L2 normalization. Because  $\varphi^S$  learns from  $\varphi^T$ , training can be unstable with large lr values if  $\varphi^T$  frequently produces outlier points based on a lack of generalization or excessive input noise. To alleviate such unwanted effects, we assign additional weight to an objective if the distance from the embedding of teacher  $e_{\alpha_k}^T$  to the closest label embedding  $\hat{d}_{\alpha_k}$  is small because a smaller distance indicates a more certain prediction. A large distance to  $\hat{d}_{\alpha_k}$  could indicate the presence of outliers, so such signals are suppressed during gradient updating. The final objective is formulated as

$$\ell^{Student} = \Sigma_{k=1}^{\kappa} \{ (\|\hat{d}_{\alpha_k}\|_2 \cdot \|e_{\alpha_k}^T\|_2)^{\beta} \times (1 - \|e_{\alpha_k}^S\|_2 \cdot \|e_{\alpha_k}^T\|_2) \},$$
(4)

where  $\beta$  is the hyperparameter to be found.  $\beta$  is set to 1.0 unless stated otherwise. Although we observed that  $\beta$  has a very small effect on final accuracy, this



Fig. 5. An example of the query inference process.

setting enables to use a large lr value at the very beginning of training phase by suppressing noise.

## 3.4 Query Inference

Given a query image x and ground-truth label y of an attribute type  $\alpha_k$ , the attribute-specific embedding of  $\alpha_k$  can be calculated as  $e_{\alpha_k}^S = \varphi_{\alpha_k}^S(x; \alpha_k)$ . Let the dictionary learned at the teacher stage  $d_{\alpha_k} = \lambda_{\alpha_k}(y)$ , where  $d_{\alpha_k} \in \{d_{\alpha_k}^1, ..., d_{\alpha_k}^N\}$  and N is the number of classes included in  $\alpha_k$ . Because  $d_{\alpha_k}$  represents the center points of the feature cluster for each ground-truth class y and  $\varphi_{\alpha_k}^S$  can reproduce the same distribution with  $\varphi_{\alpha_k}^T$ , a prediction can be obtained by finding the  $\hat{d}_{\alpha_k}$  that maximizes the cosine similarity with  $e_{\alpha_k}^S$ . An example of the query inference process for predicting a *pattern* is presented in Figure 5.

## 4 Experiment

#### 4.1 Experimental Setup

**Datasets.** All models were evaluated on several fashion-related benchmarks, namely three public datasets called iMaterialist-Fashion-2018 (iMatFashion) [13], Deepfashion Category and Attribute Prediction Benchmark (DeepFashion) [23], and DARN [16], and one private dataset called FiccY. We cropped the images using a CenterNet [44] based fashion object detector unless ground-truth bounding boxes were not provided for a dataset. We omit the detailed statistics of the public datasets because they have been well described in many related works [8, 38, 13, 34]. FiccY is a private fashion dataset collected from our database and annotated by human experts. It contains 520K high-resolution images labeled for four types of attributes (*Category, Color, Pattern, and Texture*) produced by both sellers and wearers. We intentionally include the result on FiccY in our experiments because it cab reflect a real service environment. However, the majority of experiments were performed on the DeepFashion benchmark for reproducibility.

#### 4.2 Performance on Benchmarks

Multi-Label Classification. Previous studies [6, 16, 23, 8, 38, 32, 22] on attribute prediction in the fashion domain have reported R@3 scores according to attribute types using DeepFashion, so we used the same metric for fair comparison. The results are listed in Table 1. The R@3 scores of both the teacher and student are listed for various experimental settings of  $\gamma$  and lr. The results demonstrate that with the default hyperparameter of  $\gamma = 1$ , our teacher model outperforms previous methods [8, 6, 16] that use no auxiliary supervision. We determined that the optimized values of  $\gamma$  and lr differ depending on the task, meaning these hyperparameters must be selected carefully according to the target attribute type. When training a student, we observed higher gains when a greater lr value was adopted. Adjusting the  $\beta$  value enables the use of greater lr values while avoiding instability at the beginning of training. The student trained by the tuned teachers yields very similar results compared to state-of-the-art domain-specific methods [23, 38, 32, 22].

**Table 1.** Top-k recall for attribute prediction on the DeepFashion [23] dataset. \*marked methods use additional labels for training, such as landmark, pose, or text descriptions.  $\gamma = tune$  indicates that the best-performing  $\gamma$  was selected according to the attribute types. Recall was measured when the model achieved the best F1@1 score. *Style* is not compared to the other methods because style scores were not reproducible using publicly released code. Overall  $1^{st}/2^{nd}$  best in **blue/green**.

	Cate	egory	Tex	ture	Fal	oric	Sha	аре	Pa	art	St	yle	А	.ll
Method	top-3	top-5	top-3	top-5	top-3	top-5	top-3	top-5	top-3	top-5	top-3	top-5	top-3	top-5
WTBI [6]	43.73	66.26	24.21	32.65	25.38	36.06	23.39	31.26	26.31	33.24	-	-	27.46	35.37
DARN [16]	59.48	79.58	36.15	48.15	36.64	48.52	35.89	46.93	39.17	50.14	-	-	42.35	51.95
<ul> <li>FashionNet [23]</li> </ul>	82.58	90.17	37.46	49.52	39.30	49.84	39.47	48.59	44.13	54.02	-	-	45.52	54.61
Corbiere et al. [8]	86.30	92.80	53.60	63.20	39.10	48.80	50.10	59.50	38.80	48.90	-	-	23.10	30.40
<ul> <li>* Wang et al. [38]</li> </ul>	90.99	95.78	50.31	<b>65.48</b>	40.31	48.23	53.32	61.05	40.65	56.32	-	-	51.53	60.95
* VSAM + FL [32]	-	-	56.28	65.45	41.73	52.01	55.69	<b>65.40</b>	43.20	53.95	-	-	-	-
* Liu et al. [22]	91.16	96.12	56.17	65.83	43.20	<b>53.52</b>	58.28	67.80	46.97	57.42	-	-	-	-
$Ours(T)(\gamma = 0)$	88.45	91.21	55.91	62.02	43.90	50.54	55.92	61.46	45.40	51.55	30.13	35.44	-	-
$Ours(T)(\gamma = 1)$	89.32	92.92	57.89	64.71	44.51	51.57	56.22	62.34	46.09	52.73	30.88	36.19	-	-
$Ours(T)(\gamma = tune)$	90.01	93.74	58.14	65.29	<b>44.93</b>	52.18	56.74	62.66	46.83	53.66	31.31	37.13	-	-
$Ours(S)(\gamma = 1, \beta = 1)$	89.65	93.29	58.33	65.17	<b>45.05</b>	52.26	56.91	63.17	<b>46.98</b>	54.30	32.30	38.41	57.85	<b>63.91</b>
$\underline{\mathrm{Ours}(\mathrm{S})(\gamma=tune,\beta=1)}$	90.17	93.98	57.58	64.66	44.87	52.47	57.34	63.69	47.36	54.99	<b>32.6</b> 4	39.24	58.02	64.35

LE versus Binary Cross-Entropy In this section, we analyze the strength of LE compared to the most commonly used method, namely binary cross-entropy (BCE). Table 2 compares our implementation to the existing method. The student model is trained by the teacher networks, which are trained using LE. The state-of-the-art method based on DeepFashion uses BCE for training. One can see that LE consistently yields better results when k=1, implying that prediction is more accurate with a small number of trials. We directly compared the LE and BCE methods by replacing the objective of our teacher model with BCE. The results in Table 3 reveal consistent improvements in terms of recall at k=1 for all evaluation datasets.



Fig. 6. Effects of training hyperparameters evaluated on DeepFashion: (a) gamma  $\gamma$ , (b) embedding dimension D, and (c) training data size. These plots represent analysis of the changes in R@3 (gain or degradation) compared to the baseline trained with (a)  $\gamma = 0$ , (b) D = 256, and (c) a data size equal to the size of the entire training set. Part (c) presents results for both the student (S) and teacher (T).

Hyperparameter tuning. The optimal value of  $\gamma$  for training a teacher network is analyzed in Figure 6 (a). While the use of focal loss consistently improves the results by up to 3% in terms of R@3, we found that the optimal value differs depending on the attribute types. Therefore, it is highly recommended to tune  $\gamma$  carefully for each attribute type. The feature dimension is also an important engineering factor. Figure 6 (b) presents the effects of the embedding dimension *D*. Better results can be observed for higher dimension, but processing high-dimensional features may require more resources and inference time. The result in Table 1 demonstrate that tuning the hyperparameters of not only  $\gamma$  and *D*, but also lr, significantly improves performance. The combination of optimal values is different for each attribute type.

**Training depending on data size.** Figure 6 (c) presents the performance degradation when the number of usable training images decreases. Overall R@3 scores are presented for both the teacher and student networks. The number of training images was intentionally limited to observe how performance degrades with a shortage of training data. As expected, a lack of training data significantly degrades the overall results. It is noteworthy that the student always outperforms the teachers slightly in terms of gain, even though the teachers are very weak. However, this also indicates that the final performance of the student is strongly bounded by that of teachers. We found no relationship between training data size and the amount of performance gain for a student compared to a teacher.

Table 2. The top-k F1 scores and recall values evaluated on DeepFashion.

Method	F1@1	R@1	F1@3	R@3	F1@5	R@5
Liu et al. [22]	31.35	37.17	22.50	57.30	16.82	65.66
$\overline{\text{Ours}(S)(\gamma = tune, \beta = 1)}$	33.53	40.11	22.41	58.02	16.16	64.35

11

Dataset	Category	Pattern	$\operatorname{Color}$	Texture	Avg. $\varDelta$
BCE	85.12	67.46	65.81	57.43	-
LE	85.52	67.71	66.31	57.98	+0.62%

Table 3. Comparison between LE and BCE. The F1@1 scores were measured on FiccY.

#### 4.3 Effectiveness of Distillation

Comparison between a teacher and student. In Table 4, we compare classification performances between a teacher and student on the DeepFashion, iMat-Fashion, and FiccY datasets. Results were evaluated for all three datasets to determine if the proposed method exhibits consistent experimental trends in different domains. The results illustrate that the student achieves a better score than the teachers in all cases. Our interpretation is that the use of unlabeled images can induce improvements because the negative effects caused by missed annotations are successfully suppressed during training by matching features from teacher and student for supervision. As a proof, only slight improvements in terms of F1 score are observed for *categories* which has relatively dense annotations than the others. Our MTSS approach is a semi-supervised form of multitask learning, so we analyze the benefits of learning representations simultaneously from multiple tasks in Table 5. Given four different attribute types for FiccY, we compare students trained with the supervision of either single or multiple domain experts. We found consistent improvements in both cases. However, higher improvement is achieved in general with supervision from all four teachers than a single teacher. We can conclude that the proposed model benefits from knowledge transfer between different tasks, which were represented by different attribute types in this experiment. Figure 8 analyzes the differences in R@3 scores between teachers and students for each class with  $\alpha_{\kappa} = shape$ . Although the result by class is generally improved based on distillation, performance degradation occurs in some classes. Such results are found when the recall of a teacher is poor meaning that poorly generalized teacher networks can negatively affect optimization for minor classes, while having a small effect on overall performance. Another possibility is that certain rare attributes lead to

	iMatFashion [13]										_		
			Categor	y Ger	nder M	laterial	Pattern	Sty	vle l	Veckli	ne Sleev	e Colo	r
		Т	42.50	88	.50	49.27	37.84	27.	70	40.11	76.72	47.75	5
		$\mathbf{S}$	42.43	90	.46	50.45	38.04	28.	60	41.93	3 77.66	6 49.69	9
			DeepFa	ashior	n [23]						FiccY		
	Category	Γ	exture F	abric	Shape	e Part	Style		Cate	egory	Pattern	Color	Texture
Т	74.45		25.61 2	24.97	30.45	21.12	12.63	Т	85	.10	67.69	66.15	57.74
$\mathbf{S}$	74.85		26.08 2	5.14	30.81	21.40	13.29	$\mathbf{S}$	85	.78	68.31	67.60	59.19

**Table 4.** Performance comparison between a teacher (T) and student (S) model. F1@1 scores are measured on the iMatFashion, DeepFashion, and FiccY datasets.

**Table 5.** R@3 gain of a student model compared to teachers with either single (S-Single) or multiple (S-Multi) teachers evaluated on FiccY.  $\Delta_{single}$  and  $\Delta_{multi}$  denote the results for S-Single and S-Multi, respectively, as percentages.

Attribute	Т	S-Single	$\Delta_{single}$	S-Multi	$\Delta_{multi}$	$\Delta_{multi}$ - $\Delta_{single}$
Category	96.70	97.30	+0.62%	97.30	+0.62%	+0.00%
Pattern	90.11	90.73	+0.69%	91.14	+1.14%	+0.45%
Color	88.97	89.54	+0.64%	89.98	+1.14%	+0.50%
Texture	92.06	93.57	+1.64%	94.22	+2.35%	+0.71%

underfitting for certain classes, even with a large number of training epochs. Regardless, the overall results with distillation for all classes exhibit consistent improvement.

**t-SNE visualization of distributions.** The distribution of teacher and student embeddings is visualized in Figure 7. The result illustrates that the student reproduces almost the same distribution as the teacher. We omitted visualizations for attribute types other than *patterns* because similar results were observed for all attribute types. The feature clusters for each prediction class are clearly formed for both the teacher and student.



Fig. 7. T-distributed stochastic neighbor embedding (t-SNE) [25] visualization of pattern attribute types for FiccY. The left and right parts present the distributions of image embeddings extracted from the teacher and student models, respectively.

## 4.4 Robustness and Generalization

**Cross-domain adaptability.** In this subsection, we examine the capability of SSL to transfer knowledge for adaptation to a new target domain [27]. Validating the cross-domain adaptability of VAP is difficult because the classes provided

Model	Train set	Test set	Flower	Stripe	Dot	Check	Leopard	All
Т	DeepFashion	DARN	70.46	66.11	67.47	51.00	40.90	64.31
$\mathbf{S}$	DeepFashion	DARN	71.16	67.81	68.97	52.02	43.97	65.15
$S^{\dagger}$	DeepFashion+DARN	DARN	71.84	68.70	70.03	52.62	44.09	65.49
Т	iMatFashion	DARN	74.73	75.33	74.25	68.53	51.84	72.34
$\mathbf{S}$	iMatFashion	DARN	75.37	76.74	74.90	69.77	55.72	73.10
$S^{\dagger}$	iMatFashion+DARN	DARN	75.81	76.78	75.20	70.34	57.61	73.22

**Table 6.** Cross-domain adaptability. F1@1 scores of five commonly appearing classes were measured on DARN [16]. The symbol <sup>†</sup> indicates a model trained with both original training set images and external large-scale unlabeled images.

in each dataset differ. For fair comparison, we carefully selected five commonly appearing classes (*Flower, Stripe, Dot, Check, and Leopard*) for three datasets (*DeepFashion, iMatFashion, and DARN*) and created subsets containing only images from one of the selected classes. The results in Table 6 reveal that training a student is effective for cross-domain classification. We observe clear improvements for all five classes relative to the teacher models. We also examined the strength of SSL useful for deploying a model in a realistic environment. Because the training process for a student requires no annotations, the training images for two distinct domains (DARN and DeepFashion/iMatFashion) were mixed for training. This combination provided a significant performance boost for the testing set. This could be an useful feature in real production environments because a model can be easily fine-tuned for a service domain after being trained in a constrained environment.

Robustness to corruption and perturbation. Hendrycks *et al.* [14] introduced a benchmark for image classifier robustness to corruption and perturbation. Images were transformed using 15 types of algorithmically generated corruptions with five levels of severity, resulting in a total of  $15 \times 5$  corrupted images. By following this same procedure, we created a corrupted version of the DeepFashion dataset and measured the drop in R@3 compared to the clean version. The result are listed in Table 7. Testing on the corrupted images degrades the overall R@3 score by an average of approximately 25%. In Table 6, the benefits of mixing two distinct domains in terms of domain adaptation were

**Table 7.** Results for robustness to corruption and perturbation in inputs. Recall was measured on both clean and corrupted versions of the DeepFashion dataset. The symbol † indicates a model trained with both clean images and 1M, 10M, or 15M corrupted images.

Model	Mixing	Clean R@3	Δ	Corrupted R@3	Δ
Baseline(S)	No	57.47	-	42.96	-
$S^{\dagger}$ -1M	Yes	57.75	+0.49%	<b>43.96</b>	+2.33%
$S^{\dagger}$ -10M	Yes	57.80	+0.57%	43.88	+2.14%
$S^{\dagger}$ -15M	Yes	57.74	+0.47%	43.84	+2.05%



**Fig. 8.** R@3 scores for each class of attribute type *shape* for the DeepFashion dataset. The yellow line represents the changes in performance before and after training the student. The green line is a baseline representing a change of zero percent.

explored. Following the same strategy, we mixed clean images with 1M, 10M, and 15M randomly sampled corrupted images and trained a student model to improve robustness to corruption. It should be noted that the mixed images were only used for training the student, meaning the teachers never saw the corrupted images. The results reveal that a student trained with mixed images achieves an increase of up to 2.33% in terms of R@3 score compared to a baseline student trained using only clean images. Additionally, performing fine-tuning using corrupted images does not degrade the model's performance on clean images. In fact, such fine-tuning slightly improves performance on clean images. The greatest performance gain can be observed when including only 1M corrupted images. The optimal ratio between clean and corrupted images should be identified because the noisy signals generated by corrupted images can make optimization less effective if they dominate the clean images.

# 5 Conclusion

In this paper, we proposed an MTSS approach to solving the VAP problem. Our method trains a unified model according to multiple domain experts, which enables it to predict multiple attributes that appear simultaneously in objects using a single forward operation. The core idea of MTSS is to transfer knowledge by forcing a student to reproduce the feature distributions learned by teachers. We demonstrated that such a strategy is highly effective for VAP, which suffers from a lack of effective annotations. Furthermore, our method can achieve competitive results on benchmarks using attribute labels alone, and improve robustness and domain adaptability without sacrificing accuracy fine-tuning with unlabeled images.

# References

- Abdulnabi, A.H., Wang, G., Lu, J., Jia, K.: Multi-task cnn model for attribute prediction. IEEE Transactions on Multimedia 17(11), 1949–1959 (2015)
- Adhikari, S.S., Singh, S., Rajagopal, A., Rajan, A.: Progressive fashion attribute extraction. arXiv preprint arXiv:1907.00157 (2019)
- Akata, Z., Perronnin, F., Harchaoui, Z., Schmid, C.: Label-embedding for image classification. IEEE transactions on pattern analysis and machine intelligence 38(7), 1425–1438 (2015)
- Arslan, H.S., Sirts, K., Fishel, M., Anbarjafari, G.: Multimodal sequential fashion attribute prediction. Information 10(10), 308 (2019)
- 5. Cevikalp, H., Benligiray, B., Gerek, O.N.: Semi-supervised robust deep neural networks for multi-label image classification. Pattern Recognition p. 107164 (2019)
- Chen, H., Gallagher, A., Girod, B.: Describing clothing by semantic attributes. In: European conference on computer vision. pp. 609–623. Springer (2012)
- Chen, Z.M., Wei, X.S., Wang, P., Guo, Y.: Multi-label image recognition with graph convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5177–5186 (2019)
- Corbiere, C., Ben-Younes, H., Ramé, A., Ollion, C.: Leveraging weakly annotated data for fashion image retrieval and label prediction. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2268–2274 (2017)
- Dal Pozzolo, A., Caelen, O., Bontempi, G.: When is undersampling effective in unbalanced classification tasks? In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 200–215. Springer (2015)
- van Engelen, J.E., Hoos, H.H.: A survey on semi-supervised learning. Machine Learning pp. 1–68 (2019)
- Gong, C., Chang, X., Fang, M., Yang, J.: Teaching semi-supervised classifier via generalized distillation. In: IJCAI. pp. 2156–2162 (2018)
- Guo, H., Zheng, K., Fan, X., Yu, H., Wang, S.: Visual attention consistency under image transforms for multi-label image classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 729–739 (2019)
- Guo, S., Huang, W., Zhang, X., Srikhanta, P., Cui, Y., Li, Y., Belongie, S., Adam, H., Scott, M.: The imaterialist fashion attribute dataset. arXiv preprint arXiv:1906.05750 (2019)
- 14. Hendrycks, D., Dietterich, T.: Benchmarking neural network robustness to common corruptions and perturbations. arXiv preprint arXiv:1903.12261 (2019)
- Hoffer, E., Ailon, N.: Deep metric learning using triplet network. In: International Workshop on Similarity-Based Pattern Recognition. pp. 84–92. Springer (2015)
- Huang, J., Feris, R.S., Chen, Q., Yan, S.: Cross-domain image retrieval with a dual attribute-aware ranking network. In: Proceedings of the IEEE international conference on computer vision. pp. 1062–1070 (2015)
- Kawaguchi, K., Kaelbling, L.P., Bengio, Y.: Generalization in deep learning. arXiv preprint arXiv:1710.05468 (2017)
- Khodadadeh, S., Boloni, L., Shah, M.: Unsupervised meta-learning for few-shot image classification. In: Advances in Neural Information Processing Systems. pp. 10132–10142 (2019)
- Kim, J., Kim, T., Kim, S., Yoo, C.D.: Edge-labeling graph neural network for fewshot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11–20 (2019)

- 16 M. Shin.
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- Liu, J., Lu, H.: Deep fashion analysis with feature map upsampling and landmarkdriven attention. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 0–0 (2018)
- 23. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1096–1104 (2016)
- Lu, Y., Kumar, A., Zhai, S., Cheng, Y., Javidi, T., Feris, R.: Fully-adaptive feature sharing in multi-task networks with applications in person attribute classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5334–5343 (2017)
- Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(Nov), 2579–2605 (2008)
- Neyshabur, B., Bhojanapalli, S., McAllester, D., Srebro, N.: Exploring generalization in deep learning. In: Advances in Neural Information Processing Systems. pp. 5947–5956 (2017)
- Orbes-Arteainst, M., Cardoso, J., Sørensen, L., Igel, C., Ourselin, S., Modat, M., Nielsen, M., Pai, A.: Knowledge distillation for semi-supervised domain adaptation. In: OR 2.0 Context-Aware Operating Theaters and Machine Learning in Clinical Neuroimaging, pp. 68–76. Springer (2019)
- Palatucci, M., Pomerleau, D., Hinton, G.E., Mitchell, T.M.: Zero-shot learning with semantic output codes. In: Advances in neural information processing systems. pp. 1410–1418 (2009)
- Papernot, N., Abadi, M., Erlingsson, U., Goodfellow, I., Talwar, K.: Semisupervised knowledge transfer for deep learning from private training data. arXiv preprint arXiv:1610.05755 (2016)
- Park, S., Shin, M., Ham, S., Choe, S., Kang, Y.: Study on fashion image retrieval methods for efficient fashion visual search. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
- Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3967–3976 (2019)
- 32. Quintino Ferreira, B., Costeira, J.P., Sousa, R.G., Gui, L.Y., Gomes, J.P.: Pose guided attention for multi-label fashion image classification. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 0–0 (2019)
- Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. arXiv preprint arXiv:1803.00676 (2018)
- Shin, M., Park, S., Kim, T.: Semi-supervised feature-level attribute manipulation for fashion image retrieval. In: Proceedings of the British Machine Vision Conference (2019)
- Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: Advances in neural information processing systems. pp. 4077–4087 (2017)
- Socher, R., Ganjoo, M., Manning, C.D., Ng, A.: Zero-shot learning through crossmodal transfer. In: Advances in neural information processing systems. pp. 935–943 (2013)

- 37. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: Advances in Neural Information Processing Systems. pp. 1857–1865 (2016)
- Wang, W., Xu, Y., Shen, J., Zhu, S.C.: Attentive fashion grammar network for fashion landmark detection and clothing category classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4271–4280 (2018)
- Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2840–2848 (2017)
- Xie, Q., Hovy, E., Luong, M.T., Le, Q.V.: Self-training with noisy student improves imagenet classification. arXiv preprint arXiv:1911.04252 (2019)
- 41. Yalniz, I.Z., Jégou, H., Chen, K., Paluri, M., Mahajan, D.: Billion-scale semisupervised learning for image classification. arXiv preprint arXiv:1905.00546 (2019)
- Zhang, S., Song, Z., Cao, X., Zhang, H., Zhou, J.: Task-aware attention model for clothing attribute prediction. IEEE Transactions on Circuits and Systems for Video Technology (2019)
- 43. Zheng, S., Song, Y., Leung, T., Goodfellow, I.: Improving the robustness of deep neural networks via stability training. In: Proceedings of the ieee conference on computer vision and pattern recognition. pp. 4480–4488 (2016)
- 44. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019)